arXiv:1401.0699v3 [cs.DS] 28 Apr 2014

Nonuniform Graph Partitioning with Unrelated Weights

Konstantin Makarychev
Microsoft Research

Yury Makarychev*
Toyota Technological Institute at Chicago

Abstract

We give a bi-criteria approximation algorithm for the Minimum Nonuniform Partitioning problem, recently introduced by Krauthgamer, Naor, Schwartz and Talwar (2014). In this problem, we are given a graph $G = (V, E)$ on n vertices and k numbers $\rho_1, \dots, \rho_k$. The goal is to partition the graph into k disjoint sets $P_1, \dots, P_k$ satisfying $|P_i| \leq \rho_i n$ so as to minimize the number of edges cut by the partition. Our algorithm has an approximation ratio of $O(\sqrt{\log n \log k})$ for general graphs, and an $O(1)$ approximation for graphs with excluded minors. This is an improvement upon the $O(\log n)$ algorithm of Krauthgamer, Naor, Schwartz and Talwar (2014). Our approximation ratio matches the best known ratio for the Minimum (Uniform) k-Partitioning problem.

We extend our results to the case of "unrelated weights" and to the case of "unrelated d-dimensional weights". In the former case, different vertices may have different weights and the weight of a vertex may depend on the set P_i the vertex is assigned to. In the latter case, each vertex u has a d-dimensional weight $r(u, i) = (r_1(u, i), \dots, r_d(u, i))$ if u is assigned to P_i. Each set P_i has a d-dimensional capacity $c(i) = (c_1(i), \dots, c_d(i))$. The goal is to find a partition such that $\sum_{u \in P_i} r(u, i) \leq c(i)$ coordinate-wise.

1 Introduction

We study the Minimum Nonuniform Partitioning problem, which was recently proposed by Krauthgamer, Naor, Schwartz and Talwar (2014). We are given a graph $G = (V, E)$, parameter k and k numbers (capacities) $\rho_1, \dots, \rho_k$. Our goal is to partition the graph G into k pieces (bins) $P_1, \dots, P_k$ satisfying capacity constraints $|P_i| \leq \rho_i n$ so as to minimize the number of cut edges. The problem is a generalization of the Minimum k-Partitioning problem studied by Krauthgamer, Naor, and Schwartz (2009), in which all bins have equal capacity $\rho_i = 1/k$.

The problem has many applications (see Krauthgamer et al. 2014). Consider an example in cloud computing: Imagine that we need to distribute n computational tasks – vertices of the graph – among k machines, each with capacity $\rho_i n$. Different tasks communicate with each other. The amount of communication between tasks u and v equals the weight of the edges between the corresponding vertices u and v. Our goal is to distribute tasks among k machines subject to capacity constraints so as to minimize the total amount of communication between machines.[1]

The problem is quite challenging. Krauthgamer et al. (2014) note that many existing techniques do not work for this problem. Particularly, it is not clear how to solve this problem on tree graphs[2] and consequently

*Supported by NSF CAREER award CCF-1150062 and NSF grant IIS-1302662.

[1] In this example, we need to solve a variant of the problem with edge weights.

[2] Our algorithm gives a constant factor bi-criteria approximation for trees.

how to use Räcke's (2008) tree decomposition technique. Krauthgamer et al. (2014) give an $O(\log n)$ bi-criteria approximation algorithm for the problem: the algorithm finds a partition $P_1, \ldots, P_k$ such that $|P_i| \leq O(\rho_i n)$ for every i and the number of cut edges is $O(\log n\, OPT)$. The algorithm first solves a configuration linear program and then uses a new sophisticated method to round the fractional solution.

In this paper, we present a rather simple SDP based $O(\sqrt{\log n \log k})$ bi-criteria approximation algorithm for the problem. We note that our approximation guarantee matches that of the algorithm of Krauthgamer, Naor, and Schwartz (2009) for the the Minimum k-Partitioning problem (which is a special case of Minimum Nonuniform Partitioning, see above). Our algorithm uses a technique of "orthogonal separators" developed by Chlamtac, Makarychev, and Makarychev (2006) and later used by Bansal, Feige, Krauthgamer, Makarychev, Nagarajan, Naor, and Schwartz (2011) for the Small Set Expansion problem. Using orthogonal separators, it is relatively easy to get a distribution over partitions $\{P_1, \ldots, P_k\}$ such that $\mathbb{E}[|P_i|] \leq O(\rho_i n)$ for all i and the expected number of cut edges is $O(\sqrt{\log n \log(1/\rho_{min})}\, OPT)$ where $\rho_{min} = \min_i \rho_i$. The problem is that for some i, P_i may be much larger than its expected size. The algorithm of Krauthgamer et al. (2014) solves a similar problem by first simplifying the instance and then grouping parts P_i into "mega-buckets". We propose a simpler fix: Roughly speaking, if a set P_i contains too many vertices, we remove some of these vertices and re-partition the removed vertices into k pieces again. Thus we ensure that all capacity constraints are (approximately) satisfied. It turns out that every vertex gets removed a constant number of times in expectation. Hence, the re-partitioning step increases the number of cut edges only by a constant factor. Another problem is that $1/\rho_{min}$ may be much larger than k. To deal with this problem, we transform the SDP solution (eliminating "short" vectors) and redefine thresholds ρ_i so that $1/\rho_{min}$ becomes $O(k)$.

Our technique is quite robust and allows us to solve more general versions of the problem, *Nonuniform Graph Partitioning with unrelated weights* and *Nonuniform Graph Partitioning with unrelated d-dimensional weights*.

Minimum Nonuniform Graph Partitioning with unrelated weights captures the variant of the problem where we assign vertices (tasks/jobs) to unrelated machines and the weight of a vertex (the size of the task/job) depends on the machine it is assigned to.

Definition 1.1 (Minimum Nonuniform Graph Partitioning with unrelated weights)**.** *We are given a graph $G = (V, E)$ on n vertices and a natural number $k \geq 2$. Additionally, we are given k normalized measures $\mu_1, \ldots, \mu_k$ on V (satisfying $\mu_i(V) = 1$) and k numbers $\rho_1, \ldots, \rho_k \in (0, 1)$. Our goal is to partition the graph into k pieces (bins) $P_1, \ldots, P_k$ such that $\mu_i(P_i) \leq \rho_i$ so as to minimize the number of cut edges. Some pieces P_i may be empty.*

We will only consider instances of Minimum Nonuniform Graph Partitioning that have a feasible solution. We give an $O_\varepsilon(\sqrt{\log n \log \min(1/\rho_{min}, k)})$ bi-criteria approximation algorithm for the problem.

Theorem 1.2. *For every $\varepsilon > 0$, there exists a randomized polynomial-time algorithm that given an instance of Minimum Nonuniform Graph Partitioning with unrelated weights finds a partition $P_1, \ldots, P_k$ satisfying $\mu_i(P_i) \leq 5(1+\varepsilon)\rho_i$. The expected cost of the solution is at most $D \times OPT$, where OPT is the optimal value, $D = O_\varepsilon(\sqrt{\log n \log \min(1/\rho_{min}, k)})$ and $\rho_{min} = \min_i \rho_i$. For graphs with excluded minors $D = O_\varepsilon(1)$.*

Nonuniform Graph Partitioning with unrelated d-dimensional weights further generalizes the problem. In this variant of the problem, we assume that we have d resources (e.g. CPU speed, random access memory, disk space, network). Each piece P_i has $c_j(i)$ units of resource $j \in \{1, \ldots, d\}$, and each vertex u requires $r_j(u, i)$ units of resource $j \in \{1, \ldots, d\}$ when it is assigned to piece P_i. We need to partition the graph so that capacity constraints for all resources are satisfied. The d-dimensional version of Minimum (uniform)

k-Partitioning was previously studied by Amir et al. (2014). In their problem, all $\rho_i = 1/k$ are the same, and r_j's do not depend on i.

Definition 1.3 (Minimum Nonuniform Graph Partitioning with unrelated d-dimensional weights)**.** *We are given a graph $G = (V, E)$ on n vertices. Additionally, we are given non-negative numbers $c_j(i)$ and $r_j(u, i)$ for $i \in \{1, \dots, k\}$, $j \in \{1, \dots, d\}$, $u \in V$. Our goal is to find a partition of V into $P_1, \dots, P_k$ subject to capacity constraints $\sum_{u \in V} r_j(u, i) \le c_j(i)$ for every i and j so as to minimize the number of cut edges.*

We present a bi-criteria approximation algorithm for this problem.

Theorem 1.4. *For every $\varepsilon > 0$, there exists a randomized polynomial-time algorithm that given an instance of Minimum Nonuniform Graph Partitioning with unrelated d-dimensional weights finds a partition $P_1, \dots, P_k$ satisfying*

$$\sum_{v \in V} r_j(v, i) \le 5d(1 + \varepsilon)c_j(i) \quad \text{for every } i \text{ and } j.$$

The expected cost of the solution is at most $D \times OPT$, where OPT is the optimal value, $D = O_\varepsilon(\sqrt{\log n \log k})$. For graphs with excluded minors $D = O_\varepsilon(1)$.

We note that this result is a simple corollary of Theorem 1.2 we let $\mu'_i(u) = \max_j(r_j(u, i)/c_j(i))$ and then apply our result to measures $\mu_i(u) = \mu'_i(u)/\mu'_i(V)$ (we describe the details in Appendix C).

We remark that our algorithms work if edges in the graph have arbitrary positive weights. However, for simplicity of exposition, we describe the algorithms for the setting where all edge weights are equal to one. To deal with arbitrary edge weights, we only need to change the SDP objective function.

Our paper strengthens the result of Krauthgamer et al. (2014) in two ways. First, it improves the approximation factor from $O(\log n)$ to $O(\sqrt{\log n \log k})$. Second, it studies considerably more general variants of the problem, Minimum Nonuniform Partitioning with unrelated weights and Minimum Nonuniform Partitioning with unrelated d-dimensional weights. We believe that these variants are very natural. Indeed, one of the main motivations for the Minimum Nonuniform Partitioning problem is its applications to scheduling and load balancing: in these applications, the goal is to assign tasks to machines so as to minimize the total amount of communication between different machines, subject to capacity constraints. The constraints that we study in the paper are very general and analogous to those that are often considered in the scheduling literature. We note that the method developed in Krauthgamer et al. (2014) does not handle these more general variants of the problem.

2 Algorithm

SDP Relaxation. Our relaxation for the problem is based on the SDP relaxation for the Small Set Expansion (SSE) problem of Bansal et al. (2011). We write the SSE relaxation for every cluster P_i and then add consistency constraints similar to constraints used in Unique Games. For every vertex u and index $i \in \{1, \dots, k\}$, we introduce a vector $\bar{u}_i$. In the integral solution, this vector is simply the indicator variable for the event "$u \in P_i$". It is easy to see that in the integral case, the number of cut edges equals (1). Indeed, if u and v lie in the same P_j, then $\bar{u}_i = \bar{v}_i$ for all i; if u lies in $P_{j'}$ and v lies in $P_{j''}$ (for $j' \neq j''$) then $\|\bar{u}_i - \bar{v}_i\|^2 = 1$ for $i \in \{j', j''\}$ and $\|\bar{u}_i - \bar{v}_i\|^2 = 0$ for $i \notin \{j', j''\}$. The SDP objective is to minimize (1).

We add constraint (2) saying that $\mu_i(P_i) \le \rho_i$. We further add spreading constraints (4) from Bansal et al. (2011) (see also Louis and Makarychev (2014)). The spreading constraints above are satisfied in the integral solution: If $u \notin P_i$, then $\bar{u}_i = 0$ and both sides of the inequality equal 0. If $u \in P_i$, then the left hand side equals $\mu_i(P_i)$, and the right hand side equals ρ_i.

We write standard ℓ_2^2-triangle inequalities (6) and (7). Finally, we add consistency constraints. Every vertex u must be assigned to one and only one P_i, hence constraint (5) is satisfied. We obtain the following SDP relaxation.

SDP Relaxation

$$\min \frac{1}{2} \sum_{i=1}^{k} \sum_{(u,v)\in E} \|\bar{u}_i - \bar{v}_i\|^2 \tag{1}$$

subject to

$$\sum_{u\in V} \|\bar{u}_i\|^2 \mu_i(u) \le \rho_i \qquad \text{for all } i \in [k] \tag{2}$$

$$\sum_{v\in V} \langle \bar{u}_i, \bar{v}_i \rangle \mu_i(v) \le \|\bar{u}_i\|^2 \rho_i \tag{3}$$

$$\text{for all } u \in V,\ i \in [k] \tag{4}$$

$$\sum_{i=1}^{k} \|\bar{u}_i\|^2 = 1 \qquad \text{for all } u \in V \tag{5}$$

$$\|\bar{u}_i - \bar{v}_i\|^2 + \|\bar{v}_i - \bar{w}_i\|^2 \ge \|\bar{u}_i - \bar{w}_i\|^2 \qquad \text{for all } u, v, w \in V,\ i \in [k] \tag{6}$$

$$0 \le \langle \bar{u}_i, \bar{v}_i \rangle \le \|\bar{u}_i\|^2 \qquad \text{for all } u, v \in V,\ i \in [k] \tag{7}$$

Small Set Expansion and Orthogonal Separators. Our algorithm uses a technique called "orthogonal separators". The notion of orthogonal separators was introduced in Chlamtac, Makarychev, and Makarychev (2006), where it was used in an algorithm for Unique Games. Later, Bansal et al. (2011) showed that the following holds. If the SDP solution satisfies constraints (3), (4), (6), and (7), then for every $\varepsilon \in (0,1)$, $\delta \in (0,1)$, and $i \in [k]$, there exist a distortion $D_i = O_\varepsilon(\sqrt{\log n \log(1/(\delta\rho_i))})$, and a probability distribution over subsets of V such that for a random set $S_i \subset V$ ("orthogonal separator") distributed according to this distribution, we have for $\alpha = 1/n$,

- $\mu_i(S_i) \le (1+\varepsilon)\rho_i$ (always);
- For all u, $\Pr(u \in S_i) \in [(1-\delta)\alpha\|\bar{u}_i\|^2, \alpha\|\bar{u}_i\|^2]$;
- For all $(u,v) \in E$, $\Pr(u \in S_i, v \notin S_i) \le \alpha D_i \cdot \|\bar{u}_i - \bar{v}_i\|^2$.

We let $D = \max_i D_i$. This statement was proved in Bansal et al. (2011) implicitly, so for completeness we prove it in the Appendix — see Theorem A.1. For graphs with excluded minors and bounded genus graphs, $D = O_\varepsilon(1)$.

Algorithm. Let us examine a somewhat naïve algorithm for the problem inspired by the algorithm of Bansal et al. (2011) for Small Set Expansion. We shall maintain the set of active (yet unassigned) vertices $A(t)$. Initially, all vertices are active, i.e. $A(0) = V$. At every step t, we pick a random index $i \in \{1, \dots, k\}$ and sample an orthogonal separator $S_i(t)$ as described above. We assign all active vertices from $S_i(t)$ to the bin number i:

$$P_i(t+1) = P_i(t) \cup (S_i(t) \cap A(t)),$$

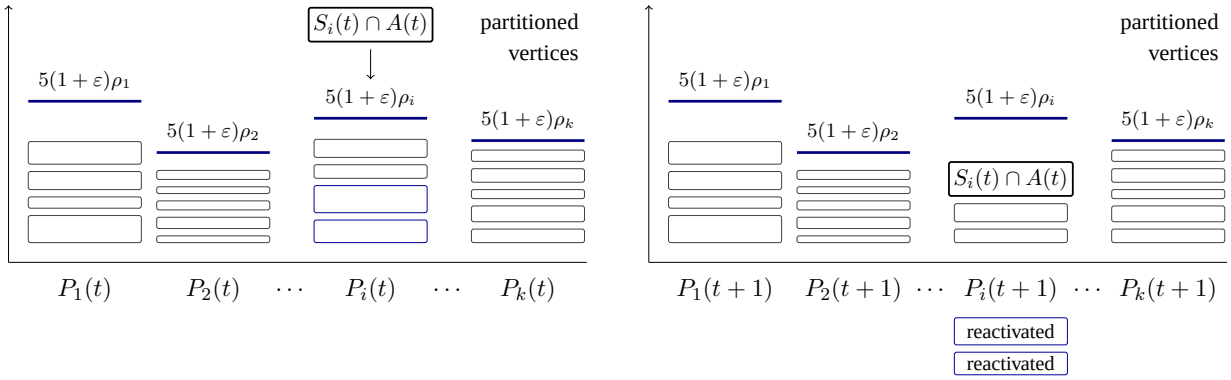

Figure 1: The figure shows how we update sets $P_i(t)$ in iteration t. In this figure, rectangles represent layers of vertices in sets $P_i(t)$ (on the left) and $P_i(t+1)$ (on the right). All vertices in these layers are inactive (they are already partitioned). Blue horizontal lines show capacity constraints. In the example shown in the figure, we add set $S_i(t) \cap A(t)$ to $P_i(t)$. The measure of the obtained set is greater than $5(1+\varepsilon)\rho_i$, and so we remove the two bottom layers from $P_i(t) \cup (S_i(t) \cap A(t))$ (the removed layers are shown in blue). We get a set of measure at most $5(1+\varepsilon)\rho_i$. Vertices in the removed layers are reactivated after the iteration is over.

and mark all newly assigned vertices as inactive i.e., we let $A(t+1) = A(t) \setminus S_i(t)$. We stop when the set of active vertices $A(t)$ is empty. We output the partition $\mathcal{P} = \{P_1(T), \ldots, P_k(T)\}$, where T is the index of the last iteration.

We can show that the number of edges cut by the algorithm is at most $O(D \times OPT)$, where D is the distortion of orthogonal separators. Furthermore, the expected weight of each P_i is $O(\rho_i)$. However, weights of some pieces may significantly deviate from the expectation and may be much larger than ρ_i. So we need to alter the algorithm to guarantee that all sizes are bounded by $O(\rho_i)$ simultaneously. We face a problem similar to the one Krauthgamer, Naor, Schwartz and Talwar (2014) had to solve in their paper. Their solution is rather complex and does not seem to work in the weighted case. Here, we propose a very simple fix for the naïve algorithm we presented above. We shall store vertices in every bin in layers. When we add new vertices to a bin at some iteration, we put them in a new layer on top of already stored vertices. Now, if the weight of the bin number i is greater than $5(1+\varepsilon)\rho_i$, we remove bottom layers from this bin so that its weight is at most $5(1+\varepsilon)\rho_i$. Then we mark the removed vertices as active and jump to the next iteration. It is clear that this algorithm always returns a solution satisfying $\mu_i(P_i) \leq 5(1+\varepsilon)\rho_i$ for all i. But now we need to prove that the algorithm terminates, and that the expected number of cut edges is still bounded by $O(D \times OPT)$.

Before proceeding to the analysis, we describe the algorithm in detail.

Algorithm for Nonuniform Partitioning with Unrelated Weights

Input: a graph $G = (V, E)$ on n vertices; a positive integer $k \leq n$; a sequence of numbers $\rho_1, \ldots, \rho_k \in (0, 1)$ (with $\rho_1 + \cdots + \rho_k \geq 1$); weights $\mu_i : V \to \mathbb{R}^+$ (with $\mu_i(V) = 1$).

Output: a partitioning of vertices into disjoint sets $P_1, \ldots, P_k$ such that $\mu_i(P_i) \leq 5(1+\varepsilon)\rho_i$.

- The algorithm maintains a partitioning of V into a set of active vertices $A(t)$ and k sets $P_1(t), \ldots P_k(t)$, which we call bins. For every inactive vertex $u \notin A(t)$, we remember its depth in the bin it belongs

to. We denote the depth by $\mathrm{depth}_u(t)$. If $u \in A(t)$, then we let $\mathrm{depth}_u(t) = \perp$.

- Initially, set $A(0) = V$; and $P_i(0) = \varnothing$, $\mathrm{depth}_u(t) = \perp$ for all i; $t = 0$.
- **while** $A(t) \neq \varnothing$
  1. Pick an index $i \in \{1, \ldots, k\}$ uniformly at random.
  2. Sample an orthogonal separator $S_i(t) \subset V$ with $\delta = \varepsilon/4$ as described in Section 2.
  3. Store all active vertices from the set $S_i(t)$ in the bin number i. If $\mu_i(P_i(t) \cup (S_i(t) \cap A(t))) \leq 5(1+\varepsilon)\rho_i$, then simply add these vertices to $P_i(t+1)$:
  $$P_i(t+1) = P_i(t) \cup (S_i(t) \cap A(t)).$$
  Otherwise, find the largest depth d such that $\mu_i(P_i(t+1)) \leq 5(1+\varepsilon)\rho_i$, where
  $$P_i(t+1) = \{u \in P_i(t) : \mathrm{depth}_u(t) \leq d\} \cup (S_i(t) \cap A(t)).$$
  In other words, add to the bin number i vertices from $S_i(t) \cap A(t)$ and remove vertices from the bottom layers so that the weight of the bin is at most $5(1+\varepsilon)\rho_i$.
  4. If we put at least one new vertex in the bin i at the current iteration, that is, if $A(t) \cap S_i(t) \neq \varnothing$, then set the depth of all newly stored vertices to 1; increase the depth of all other vertices in the bin i by 1.
  5. Update the set of active vertices: let $A(t+1) = V \setminus \bigcup_j P_j(t+1)$ and $\mathrm{depth}_u(t+1) = \perp$ for $u \in A(t+1)$. Let $t = t+1$.
- Set $T = t$ and return the partitioning $P_1(T), \ldots, P_k(T)$.

Note that Step 3 is well defined. We can always find an index d such that $\mu_i(P_i(t+1)) \leq 5(1+\varepsilon)\rho_i$, because for $d = 0$, we have $P_i(t+1) = S_i(t) \cap A(t)$ and thus

$$\mu(P_i(t+1)) = \mu_i(S_i(t) \cap A(t)) \leq \mu_i(S_i(t)) \leq (1+\varepsilon)\rho_i < 5(1+\varepsilon)\rho_i,$$

by the first property of orthogonal separators.

Analysis. We will first prove Theorem 2.1 that states that the algorithm has approximation factor $D = O_\varepsilon(\sqrt{\log n \log(1/\rho_{min})})$ on arbitrary graphs, and $D = O_\varepsilon(1)$ on graphs excluding a minor. Then we will show how to obtain $D = O_\varepsilon(\sqrt{\log n \log k})$ approximation on arbitrary graphs (see Appendix B). To this end, we will transform the SDP solution and redefine measures μ_i and capacities ρ_i so that $\rho_{min} \geq \delta/k$, then apply Theorem 2.1. The new SDP solution will satisfy all SDP constraints except possibly for constraint (5); it will however satisfy a relaxed constraint

$$\sum_{i=1}^{k} \|\bar{u}_i\|^2 \in [1-\delta, 1] \qquad \text{for all } u \in V. \tag{5$'$}$$

Thus in Theorem 2.1, we will assume only that the solution satisfies the SDP relaxation with constraint (5) replaced by constraint (5$'$).

Theorem 2.1. *The algorithm returns a partitioning $P_1(T), \dots, P_k(T)$ satisfying $\mu_i(P_i) \le 5(1+\varepsilon)\rho_i$. The expected number of iterations of the algorithm is at most $\mathbb{E}[T] \le 4n^2k + 1$ and the expected number of cut edges is at most $O(D \times SDP) = O(D \times OPT)$, where $D = O_\varepsilon(\sqrt{\log n \log(1/\rho_{min})})$ is the distortion of orthogonal separators; $\rho_{min} = \min_i \rho_i$. If the graph has an excluded minor, then $D = O_\varepsilon(1)$ (the constant depends on the excluded minor).*

We assume only that the SDP solution given to the algorithm satisfies the SDP relaxation with constraint (5) replaced by constraint (5′).

As we mentioned earlier, the algorithm always returns a valid partitioning. We need to verify that the algorithm terminates in expected polynomial time, and that it produces cuts of cost at most $O(D \times OPT)$ (see also Remark C.1).

The state of the algorithm at iteration t is determined by the sets $A(t)$, $P_1(t), \dots, P_k(t)$ and the depths of the elements. We denote the state by $\mathcal{C}(t) = \{A(t), P_1(t), \dots, P_k(t), \mathrm{depth}(t)\}$. Observe that the probability that the algorithm is in the state $\mathcal{C}^*$ at iteration $(t+1)$ is determined only by the state of the algorithm at iteration t. It does not depend on t (given $\mathcal{C}(t)$). So the states of the algorithm form a Markov random chain. The number of possible states is finite (since the depth of every vertex is bounded by n). To simplify the notation, we assume that for $t \ge T$, $\mathcal{C}(t) = \mathcal{C}(T)$. This is consistent with the definition of the algorithm — if we did not stop the algorithm at time T, it would simply idle, since $A(t) = \varnothing$, and thus $S_i(t) \cap A(t) = \varnothing$ for $t \ge T$.

We are interested in the probability that an inactive vertex u which lies *in the top layer* of one of the bins (i.e., $u \notin A(t)$ and $\mathrm{depth}_u(t) = 1$) is removed from that bin within m iterations. We let

$$f(m, u, \mathcal{C}^*) = \Pr(\exists t \in [t_0, t_0+m] \text{ s.t. } u \in A(t) \mid \mathcal{C}(t_0) = \mathcal{C}^*, \mathrm{depth}_u(t_0) = 1).$$

That is, $f(m, u, \mathcal{C}^*)$ is the probability that u is removed from the bin i at one of the iterations $t \in [t_0, t_0+m]$ given that at iteration t_0 the state of the algorithm is $\mathcal{C}^*$ and u is in the top layer of the bin i. Note that the probability above does not depend on t_0 and thus $f(m, u, \mathcal{C}^*)$ is well defined. We let

$$f(m) = \max_{u \in V} \max_{\mathcal{C}^*} f(m, u, \mathcal{C}^*).$$

Our fist lemma gives a bound on the expected number of steps on which a vertex u is active in terms of $f(m)$.

Lemma 2.2. *For every possible state of the algorithm $\mathcal{C}^*$, every vertex u, and natural number t_0,*

$$\sum_{t=t_0}^{t_0+m} \Pr(u \in A(t) \mid \mathcal{C}(t_0) = \mathcal{C}^*) \le \frac{k}{(1-2\delta)\alpha(1-f(m-1))}. \tag{8}$$

Proof. The left hand side of inequality (8) equals expected number (conditioned on $\mathcal{C}(t_0) = \mathcal{C}^*$) of iterations t in the interval $[t_0, t_0+m]$ at which u is active i.e., $u \in A(t)$. Our goal is to upper bound this quantity.

Initially, at time t_0, u is active or inactive. At every time t when u is active, u is thrown in one of the bins P_i with probability at least (here, we use that the SDP solution satisfies constraint (5′))

$$\frac{1}{k}\sum_{i=1}^{k}(1-\delta)\alpha\|\bar{u}_i\|^2 \ge \frac{(1-2\delta)\alpha}{k}.$$

So the expected number of iterations passed since u becomes active till u is stored in one of the bins and thus becomes inactive is at most $k/((1-2\delta)\alpha)$.

Suppose that u is stored in a bin i at iteration t, then $u \in P_i(t+1)$ and $\mathrm{depth}_u(t+1) = 1$. Thus, the probability that u is reactivated till iteration t_0+m i.e., the probability that for some $\tau \in [(t+1), t_0+m] \subset [(t+1), (t+1)+(m-1)]$, $u \in A(\tau)$ is at most $f(m-1)$. Consequently, the expected number of iterations $t \in [t_0, t_0+m]$ at which u is active is bounded by

$$\frac{k \cdot 1}{(1-2\delta)\alpha} + \frac{k \cdot f(m)}{(1-2\delta)\alpha} + \frac{k \cdot f^2(m)}{(1-2\delta)\alpha} + \cdots = \frac{k}{(1-2\delta)\alpha(1-f(m))}.$$

□

We now show that $f(m) \le 1/2$ for all m.

Lemma 2.3. *For all natural m, $f(m) \le 1/2$.*

Proof. We prove this lemma by induction on m. For $m = 0$, the statement is trivial as $f(0) = 0$.

Consider an arbitrary state $\mathcal{C}^*$, bin i^*, vertex u, and iteration t_0. Suppose that $\mathcal{C}(t_0) = \mathcal{C}^*$, $u \in P_{i^*}(t_0)$ and $\mathrm{depth}_u(t_0) = 1$ i.e., u lies in the top layer in the bin i^*. We need to estimate the probability that u is removed from the bin i^* till iteration $t_0 + m$. The vertex u is removed from the bin i^* if and only if at some iteration $t \in \{t_0, \ldots, t_0+m-1\}$, u is "pushed away" from the bin by new vertices (see Step 2 of the algorithm). This happens only if the weight of vertices added to the bin i^* at iterations $\{t_0, \ldots, t_0+m-1\}$ plus the weight of vertices in the first layer of the bin at iteration t_0 exceeds $5(1+\varepsilon)\rho_i$. Since the weight of vertices in the first layer is at most $(1+\varepsilon)\rho_i$, the weight of vertices added to the bin i^* at iterations $\{t_0, \ldots, t_0+m-1\}$ must be greater than $4(1+\varepsilon)\rho_{i^*}$.

We compute the expected weight of vertices thrown in the bin i^* at iterations $t \in \{t_0, \ldots, t_0+m-1\}$. Let us introduce some notation: $M = \{t_0, \ldots, t_0+m-1\}$; $i(t)$ is the index i chosen by the algorithm at the iteration t. Let X_{M,i^*} be the weight of vertices thrown in the bin i^* at iterations $t \in M$. Then,

$$\begin{aligned}\mathbb{E}\big[X_{M,i^*} \mid \mathcal{C}(t_0) = \mathcal{C}^*\big] &= \mathbb{E}\Big[\sum_{\substack{t \in M \\ s.t.\ i(t)=i^*}} \mu_{i^*}\big(S_{i^*}(t) \cap A(t)\big) \mid \mathcal{C}(t_0) = \mathcal{C}^*\Big] \qquad (9)\\ &= \sum_{t \in M} \sum_{v \in V} \Pr\big(i(t) = i^* \text{ and } v \in S_{i^*}(t) \cap A(t) \;\mid\; \mathcal{C}(t_0) = \mathcal{C}^*\big)\mu_{i^*}(v).\end{aligned}$$

The event "$i(t) = i^*$ and $v \in S_{i^*}(t)$" is independent from the event "$v \in A(t)$ and $\mathcal{C}(t_0) = \mathcal{C}^*$". Thus,

$$\begin{aligned}&\Pr\big(i(t) = i^* \text{ and } v \in S_{i^*}(t) \cap A(t) \;\mid\; \mathcal{C}(t_0) = \mathcal{C}^*\big)\\ &\qquad = \Pr\big(i(t) = i^* \text{ and } v \in S_{i^*}(t)\big) \cdot \Pr\big(v \in A(t) \;\mid\; \mathcal{C}(t_0) = \mathcal{C}^*\big).\end{aligned}$$

Since $i(t)$ is chosen uniformly at random in $\{1, \ldots, k\}$, we have $\Pr(i(t) = i^*) = 1/k$. Then, by property 2 of orthogonal separators, $\Pr(v \in S_{i^*}(t) \;\mid\; i(t) = i^*) \le \alpha\|\bar{v}_{i^*}\|^2$. We get

$$\Pr\big(i(t) = i^* \text{ and } v \in S_{i^*}(t) \cap A(t) \;\mid\; \mathcal{C}(t_0) = \mathcal{C}^*\big) \le \frac{\alpha\|\bar{v}_{i^*}\|^2}{k} \cdot \Pr\big(v \in A(t) \;\mid\; \mathcal{C}(t_0) = \mathcal{C}^*\big).$$

We now plug this expression in (9) and use Lemma 2.2,

$$\begin{aligned}\mathbb{E}[X_{M,i^*} \mid \mathcal{C}(t_0) = \mathcal{C}^*] &\le \sum_{v \in V} \frac{\alpha \|\bar{v}_{i^*}\|^2 \mu_{i^*}(v)}{k} \cdot \sum_{t \in M} \Pr\left(v \in A(t) \;\mid\; \mathcal{C}(t_0) = \mathcal{C}^*\right) \\ &\le \sum_{v \in V} \frac{\alpha \|\bar{v}_{i^*}\|^2 \mu_{i^*}(v)}{k} \cdot \frac{k}{(1-2\delta)\alpha(1 - f(m-1))} \\ &= \sum_{v \in V} \frac{\|\bar{v}_{i^*}\|^2 \mu_{i^*}(v)}{(1-2\delta)(1-f(m-1))}.\end{aligned}$$

Finally, observe that $1 - f(m-1) \ge 1/2$ by the inductive hypothesis, and $\sum_{v \in V} \|\bar{v}_{i^*}\|^2 \mu_{i^*}(v) \le \rho_{i^*}$ by the SDP constraint (2). Hence, $\mathbb{E}[X_{M,i^*} \mid \mathcal{C}(t_0) = \mathcal{C}^*] \le 2\rho_{i^*}/(1-2\delta)$. By Markov's inequality,

$$\Pr\left(X_{M,i^*} \ge 4(1+\varepsilon)\rho_{i^*}\right) \le \frac{2\rho_{i^*}}{4(1-2\delta)(1+\varepsilon)\rho_{i^*}} \le \frac{1}{2},$$

since $\delta = \varepsilon/4$. This concludes the proof. □

As an immediate corollary of Lemmas 2.2 and 2.3, we get that for all $u \in V$,

$$\sum_{t=0}^{\infty} \Pr(u \in A(t)) = \lim_{m \to \infty} \sum_{t=0}^{m} \Pr(u \in A(t)) \le \frac{2k}{(1-2\delta)\alpha} \le \frac{4k}{\alpha}. \tag{10}$$

Proof of Theorem 2.1. We now prove Theorem 2.1. We first bound the expected running time. At every iteration of the algorithm $t < T$, the set $A(t)$ is not empty. Hence, using (10), we get

$$\mathbb{E}[T] \le \mathbb{E}\Big[\sum_{t=0}^{\infty} |A(t)|\Big] + 1 = \sum_{v \in V} \sum_{t=0}^{\infty} \Pr(v \in A(t)) + 1 \le n \cdot \frac{4k}{\alpha} + 1 = 4kn^2 + 1.$$

We now upper bound the expected size of the cut. For every edge $(u, v) \in E$ we estimate the probability that (u, v) is cut. Suppose that (u, v) is cut. Then, u and v belong to distinct sets $P_i(T)$. Consider the iteration t at which u and v are separated the first time. A priori, there are two possible cases:

1. At iteration t, u and v are active, but only one of the vertices u or v is added to some set $P_i(t+1)$; the other vertex remains in the set $A(t+1)$.

2. At iteration t, u and v are in some set $P_i(t)$, but only one of the vertices u or v is removed from the set $P_i(t+1)$.

It is easy to see that, in fact, the second case is not possible, since if u and v were never separated before iteration t, then u and v must have the same depth (i.e., $\mathrm{depth}_u(t) = \mathrm{depth}_v(t)$) and thus u and v may be removed from the bin i only together.

Consider the first case, and assume that $u \in P_{i(t)}(t+1)$ and $v \in A(t+1)$. Here, as in the proof of Lemma 2.3, we denote the index i chosen at iteration t by $i(t)$. Since $u \in P_{i(t)}(t+1)$ and $v \in A(t+1)$, we have $u \in S_{i(t)}(t)$ and $v \notin S_{i(t)}(t)$. Write

$$\begin{aligned}&\Pr(u, v \in A(t);\, u \in S_{i(t)}(t);\; v \notin S_{i(t)}(t)) = \\ &\qquad = \Pr(u, v \in A(t)) \cdot \Pr(u \in S_{i(t)}(t);\; v \notin S_{i(t)}(t)) \\ &\qquad = \Pr(u, v \in A(t)) \cdot \sum_{i=1}^{k} \frac{\Pr(u \in S_i(t);\; v \notin S_i(t) \mid i(t) = i)}{k}.\end{aligned}$$

We replace $\Pr(u, v \in A(t))$ with $\Pr(u \in A(t)) \geq \Pr(u, v \in A(t))$, and then use the inequality $\Pr(u \in S_i(t);\ v \notin S_i(t)) \leq \alpha D \ \|\bar{u}_i - \bar{v}_i\|^2$, which follows from the third property of orthogonal separators. We get

$$\Pr(u, v \in A(t);\ u \in S_{i(t)}(t);\ v \notin S_{i(t)}(t)) \leq \Pr(u \in A(t)) \times \Big(\frac{1}{k}\sum_{i=1}^{k} \alpha D \ \|\bar{u}_i - \bar{v}_i\|^2\Big).$$

Thus, the probability that u and v are separated at iteration t is upper bounded by $\Big(\Pr(u \in A(t)) + \Pr(v \in A(t))\Big) \times \Big(\frac{1}{k}\sum_{i=1}^{k} \alpha D \ \|\bar{u}_i - \bar{v}_i\|^2\Big)$. The probability that the edge (u, v) is cut (at some iteration) is at most

$$\Big(\sum_{t=0}^{\infty} \Pr(u \in A(t)) + \Pr(v \in A(t))\Big) \times \Big(\frac{1}{k}\sum_{i=1}^{k} \alpha D \ \|\bar{u}_i - \bar{v}_i\|^2\Big) \leq$$

$$\leq \frac{8k}{\alpha}\Big(\frac{1}{k}\sum_{i=1}^{k} \alpha D \ \|\bar{u}_i - \bar{v}_i\|^2\Big) = 8\sum_{i=1}^{k} D \ \|\bar{u}_i - \bar{v}_i\|^2.$$

To bound the first term on the left hand side we used inequality (10). We get the desired bound on the expected number of cut edges:

$$\sum_{(u,v)\in E} \Pr((u, v) \text{ is cut}) \leq 8 \sum_{(u,v)\in E} \sum_{i=1}^{k} D \ \|\bar{u}_i - \bar{v}_i\|^2 = 16D \cdot SDP,$$

where SDP is the SDP value. □

References

A. Amir, J. Ficler, R. Krauthgamer, L. Roditty, and O. Sar Shalom. Multiply Balanced k-Partitioning. LATIN 2014.

N. Bansal, U. Feige, R. Krauthgamer, K. Makarychev, V. Nagarajan, J. Naor, and R. Schwartz. Min-max Graph Partitioning and Small Set Expansion. FOCS 2011.

E. Chlamtac, K. Makarychev, and Y. Makarychev. How to Play Unique Games Using Embeddings. FOCS 2006.

R. Krauthgamer, J. Naor, R. Schwartz. Partitioning graphs into balanced components. SODA 2009.

R. Krauthgamer, J. Naor, R. Schwartz and K. Talwar (2014). Non-Uniform Graph Partitioning. SODA 2014.

A. Louis and K. Makarychev. Approximation Algorithm for Sparsest k-Partitioning. SODA 2014.

H. Räcke. Optimal hierarchical decompositions for congestion minimization in networks. STOC 2008.

A Orthogonal Separators

For completeness, we prove Theorem A.1.

Theorem A.1 (Bansal et al. (2011))**.** *There exists a polynomial-time algorithm that given a graph $G = (V, E)$, a measure μ on V ($\mu(V) = 1$), parameters $\rho, \varepsilon, \delta \in (0, 1)$ and a collection of vectors $\bar{u}$ satisfying the following constraints:*

$$\sum_{u \in V} \|\bar{u}\|^2 \mu(u) \le \rho \qquad \text{for all } i \in [k] \tag{11}$$

$$\sum_{v \in V} \langle \bar{u}, \bar{v} \rangle \mu(v) \le \|\bar{u}\|^2 \rho \qquad \text{for all } u \in V,\ i \in [k] \tag{12}$$

$$\|\bar{u} - \bar{v}\|^2 + \|\bar{v} - \bar{w}\|^2 \ge \|\bar{u} - \bar{w}\|^2 \qquad \text{for all } u, v, w \in V,\ i \in [k] \tag{13}$$

$$0 \le \langle \bar{u}, \bar{v} \rangle \le \|\bar{u}\|^2 \qquad \text{for all } u, v \in V,\ i \in [k] \tag{14}$$

$$\|\bar{u}\|^2 \le 1 \qquad \text{for all } u \in V,\ i \in [k] \tag{15}$$

outputs a random set $S \subset V$ ("orthogonal separator") such that

1. *$\mu(S) \le (1 + \varepsilon)\rho$ (always);*
2. *For all u, $\Pr(u \in S) \in [(1 - \delta)\alpha\|\bar{u}\|^2, \alpha\|\bar{u}\|^2]$;*
3. *For all $(u, v) \in E$, $\Pr(u \in S, v \notin S) \le \alpha D \cdot \|\bar{u} - \bar{v}\|^2$.*

Where the probability scale $\alpha = 1/n$, and the distortion $D \le O_\varepsilon(\sqrt{\log n \log(1/(\rho\delta))})$. For graphs with excluded minors, $D = O_\varepsilon(1)$.

In Chlamtac, Makarychev, and Makarychev (2006), we showed that there exists a randomized polynomial-time algorithm that outputs a random set S with the following properties (see also Bansal et al. (2011) and Louis and Makarychev (2014)):

- For all $u \in V$, $\Pr(u \in S) = \alpha \|\bar{u}\|^2$.
- For all $u, v \in V$ with $\|\bar{u} - \bar{v}\|^2 \ge \beta \min(\|\bar{u}\|^2, \|\bar{v}\|^2)$,

$$\Pr(u \in S \text{ and } v \in S) \le \frac{\alpha \min(\|\bar{u}\|^2, \|\bar{v}\|^2)}{m}.$$

- For all $(u, v) \in E$,

$$\Pr(u \in S \text{ and } v \notin S) \le \alpha D \times \|\bar{u} - \bar{v}\|^2.$$

Here $m > 0$ is a parameter of the algorithm; $\alpha = 1/n$ is a probability scale; $D \le O_\beta(\sqrt{\log n \log m})$ is the distortion. Bansal et al. (2011) showed that for graphs with excluded minors, $D = O(1)$.

Our algorithm samples S as above (with $m = 2/(\delta\varepsilon\rho)$, $\beta = \varepsilon/4$) and outputs $S' = S$ if $\mu(S) \le (1+\varepsilon)\rho$, and $S' = \varnothing$, otherwise. It is clear that $\mu(S') \le (1+\varepsilon)\rho$ (always), and thus the first property in Theorem A.1 is satisfied. Then, for $(u, v) \in E$,

$$\Pr(u \in S' \text{ and } v \notin S') \le \Pr(u \in S \text{ and } v \notin S) \le \alpha D \times \|\bar{u} - \bar{v}\|^2,$$

where $D = O_\beta(\sqrt{\log n \log m}) = O_\varepsilon(\sqrt{\log n \log(1/(\rho\delta))})$.

For every $u \in V$,

$$\Pr(u \in S') \le \Pr(u \in S) = \alpha\|\bar{u}\|^2.$$

So we only need to verify that $\Pr(u \in S') \ge \alpha(1-\delta)\|\bar{u}\|^2$. We assume $\|\bar{u}\|^2 \neq 0$. We have

$$\Pr(u \in S') = \Pr(u \in S' \mid u \in S) \cdot \Pr(u \in S) = \Pr\big(\mu(S) \le (1+\varepsilon)\rho \mid u \in S\big) \cdot \alpha\|\bar{u}\|^2.$$

We split V into two sets $A_u = \{v : \|\bar{u} - \bar{v}\|^2 \ge \beta\|\bar{u}\|^2\}$ and $B_u = \{v : \|\bar{u} - \bar{v}\|^2 < \beta\|\bar{u}\|^2\}$. We show below (see Lemma A.2) that $\mu(B_u) \le (1+\varepsilon/2)\rho$. Then,

$$\mu(S) = \mu(S \cap A_u) + \mu(S \cap B_u) \le (1+\varepsilon/2)\rho + \mu(S \cap A_u)$$

and

$$\Pr(u \in S') \ge \alpha\|\bar{u}\|^2 \cdot \Pr\big(\mu(S \cap A_u) \le \varepsilon\rho/2 \mid u \in S\big). \tag{16}$$

We estimate $\Pr\big(\mu(S \cap A_u) \ge \varepsilon\rho/2 \mid u \in S\big)$. For every $v \in A_u$, $\|\bar{u} - \bar{v}\|^2 \ge \beta\|\bar{u}\|^2$. Thus, for $v \in A_u$, $\Pr(u \in S; v \in S) \le \alpha\|\bar{u}\|^2/m$, and

$$\Pr(v \in S \mid u \in S) = \frac{\Pr(u \in S, v \in S)}{\Pr(u \in S)} \le \frac{1}{m}.$$

Therefore, $\mathbb{E}[\mu(S \cap A_u) \mid u \in S] \le \mu(A_u)/m \le 1/m$, and, by Markov's inequality,

$$\Pr\big(\mu(S \cap A_u) \ge \varepsilon\rho/2 \mid u \in S\big) \le \frac{\mathbb{E}[\mu(S) \mid u \in S]}{\varepsilon\rho/2} \le \frac{2}{m\varepsilon\rho} \le \delta.$$

We plug this bound in (16) and get the desired bound,

$$\Pr(u \in S') \ge \alpha\|\bar{u}\|^2 \cdot \Pr\big(\mu(S \cap A_u) \le \varepsilon\rho/2 \mid u \in S\big) \ge \alpha\|\bar{u}\|^2 \cdot (1-\delta).$$

We now prove Lemma A.2.

Lemma A.2. *For every $u \in S$ with $\|\bar{u}\|^2 \neq 0$, $\mu(B_u) \le (1+\varepsilon/2)\rho$.*

Proof. If $v \in B_u$, then by the definition of B_u and by inequality (14), we have

$$\|\bar{u}\|^2 - \langle \bar{u}, \bar{v}\rangle = \|\bar{u} - \bar{v}\|^2 - (\|\bar{v}\|^2 - \langle \bar{u}, \bar{v}\rangle) \le \beta\|\bar{u}\|^2.$$

Thus, $\langle \bar{u}, \bar{v}\rangle \ge (1-\beta)\|\bar{u}\|^2$. Now, we use constraint (12),

$$\begin{aligned}\mu(B_u) = \sum_{v \in B_u} \mu(v) &\le \sum_{v \in B_u} \mu(v) \cdot \frac{\langle \bar{u}, \bar{v}\rangle}{(1-\beta)\|\bar{u}\|^2} \le \frac{1}{(1-\beta)\|\bar{u}\|^2} \sum_{v \in V} \langle \bar{u}, \bar{v}\rangle \mu(v) \\ &\le \frac{(1+2\beta)}{\|\bar{u}\|^2} \cdot \rho\|\bar{u}\|^2 = (1+2\beta)\rho = (1+\varepsilon/2)\rho.\end{aligned}$$

□

B $O(\sqrt{\log n \log k})$ approximation

Theorem B.1. *There is a polynomial-time randomized algorithm that returns a partitioning $P_1(T), \dots, P_k(T)$ satisfying $\mu_i(P_i) \le 5(1+\varepsilon)\rho_i$ such that the expected number of cut edges is at most $O(D \times OPT)$, where $D = O_\varepsilon(\sqrt{\log n \log k})$.*

Proof. We perform three steps. First we solve the SDP relaxation, then transform its solution and change measures μ_i, and finally apply Theorem 2.1 to the obtained SDP solution.

We start with describing how we transform the solution. We set $\delta = \varepsilon/4$ as before. Then we choose a threshold θ uniformly at random from $[\delta/2, \delta]$. We let $\tilde{u}_i = \bar{u}_i$ if $\|\bar{u}_i\|^2 \ge \theta/k$ and $\tilde{u}_i = 0$, otherwise. It is immediate that the solution $\tilde{u}_i$ satisfies all SDP constraints except possibly constraint (5). Note, however, that it satisfies constraint (5′):

$$\sum_{i=1}^{k} \|\tilde{u}_i\|^2 = \sum_{i=1}^{k} \|\bar{u}_i\|^2 - \sum_{i:\|\tilde{u}_i\|^2<\theta/k} \|\tilde{u}_i\|^2 = 1 - \sum_{i:\|\tilde{u}_i\|^2<\theta/k} \|\tilde{u}_i\|^2 \in [1-\delta, 1].$$

Consider two vertices u and v. Assume without loss of generality that $\|\bar{u}_i\|^2 \le \|\bar{v}_i\|^2$. If either $\|\bar{u}_i\|^2 \le \|\bar{v}_i\|^2 < \theta/k$ or $\theta/k \le \|\bar{u}_i\|^2 \le \|\bar{v}_i\|^2$, then we have $\|\tilde{u}_i - \tilde{v}_i\| = \|\bar{u}_i - \bar{v}_i\|$. Otherwise, if $\|\bar{u}_i\|^2 < \theta/k \le \|\bar{v}_i\|^2$, we have

$$\|\tilde{u}_i - \tilde{v}_i\| = \|\bar{v}_i\|^2 \le \|\bar{u}_i - \bar{v}_i\|^2 + \|\bar{u}_i\|^2 = \|\bar{u}_i - \bar{v}_i\|^2 + \delta/k.$$

Therefore,

$$\mathbb{E}[\|\tilde{u}_i - \tilde{v}_i\|^2] \le \|\bar{u}_i - \bar{v}_i\|^2 + (\delta/k)\Pr\left(\|\bar{u}_i\|^2 < \theta/k \le \|\bar{v}_i\|^2\right).$$

To upper bound $\Pr\left(\|\bar{u}_i\|^2 < \theta/k \le \|\bar{v}_i\|^2\right)$, note that the random variable θ is distributed uniformly on $(\delta/2, \delta)$, so its probability density is bounded from above by $2/\delta$. We get from SDP constraint (7) that $\|\bar{v}_i\|^2 - \|\bar{u}_i\|^2 \le \|\bar{u}_i - \bar{v}_i\|^2$. Thus,

$$\Pr\left(\|\bar{u}_i\|^2 < \theta/k \le \|\bar{v}_i\|^2\right) \le (2k/\delta) \cdot \|\bar{u}_i - \bar{v}_i\|^2,$$

We have,

$$\mathbb{E}[\|\tilde{u}_i - \tilde{v}_i\|^2] \le \|\bar{u}_i - \bar{v}_i\|^2 + (\delta/k) \cdot (2k/\delta) \cdot \|\bar{u}_i - \bar{v}_i\|^2 = 3\|\bar{u}_i - \bar{v}_i\|^2.$$

We conclude that the SDP value of solution $\tilde{u}_i$ is at most $3SDP \le 3OPT$ in expectation.

Now we modify measures μ_i and capacities c_i. Let $A_i = \{u : \bar{u}_i \ne 0\}$. Define

$$\mu_i'(Z) = \mu_i(Z \cap A_i)/\mu_i(A_i) \text{ for } Z \subseteq V,$$
$$\tilde{\rho}_i = \rho_i/\mu_i(A_i)$$

(if $\mu_i(A_i) = 0$ we let $\tilde{\mu}_i = \mu_i$ and $\tilde{\rho}_i = 1$, essentially removing the capacity constraint for P_i). We have $\tilde{\mu}_i(V) = \mu_i(A_i)/\mu_i(A_i) = 1$. By (2), we get

$$\rho_i \ge \sum_{u \in V} \|\bar{u}_i\|^2 \mu_i(u) \ge \sum_{u \in A_i} \|\bar{u}_i\|^2 \mu_i(u) \ge \sum_{u \in A_i} \frac{\delta}{2k} \cdot \mu_i(u) = \frac{\delta \mu_i(A)}{2k}.$$

Therefore, $\tilde{\rho}_i = \rho_i/\mu_i(A) \ge \delta/(2k)$, and $\tilde{\rho}_{min} = \min \tilde{\rho}_i \ge \delta/(2k)$ (if $\mu_i(A_i) = 0$ then $\tilde{\rho}_i = 1 > \delta/(2k)$).

Note that since each ρ_i increases by a factor of $1/\mu_i(A_i)$ and each $\mu_i(u)$ increases by a factor at most $1/\mu_i(A_i)$, vectors $\tilde{u}_i$ satisfy SDP constraints (2) and (4), in which μ_i and ρ_i are replaced with $\tilde{\mu}_i$ and

$\tilde{\rho}_i$, respectively (assuming that $\mu_i(A_i) \neq 0$; if $\mu_i(A_i) = 0$, the constraints clearly hold). We run the algorithm from Theorem 2.1 on vectors $\tilde{u}_i$ with measures $\tilde{\mu}_i$ and capacities ρ_i. The algorithm finds a partition $P_1, \ldots, P_k$ that cuts at most $D \cdot SDP \leq D \cdot OPT$ edges, where $D = O_\varepsilon(\sqrt{\log n \log(1/\tilde{\rho}_{min})}) = O_\varepsilon(\sqrt{\log n \log k})$. We verify that the weight of each set P_i is $O(\rho_i)$. Note that $P_i \subset A_i$ since for $u \notin A_i$, $\|\tilde{u}_i\|^2 = 0$, and thus the algorithm does not add u to P_i. We have,

$$\mu_i(P_i) = \mu'_i(P_i \cap A_i) \cdot \mu_i(A_i) = \mu'_i(P_i) \cdot \mu_i(A_i) \leq 5(1+\varepsilon)\tilde{\rho}_i \cdot \mu_i(A_i) \leq 5(1+\varepsilon)\rho_i.$$

□

C Partitioning with d-Dimensional Weights

We describe how Minimum Nonuniform Graph Partitioning with unrelated d-dimensional weights reduces to Minimum Nonuniform Graph Partitioning with unrelated weights. Consider an instance $\mathcal{I}$ of Minimum Nonuniform Graph Partitioning with unrelated d-dimensional weights. Let $\mu'_i(u) = \max_j(r_j(u,i)/c_j(i))$. Then define measures $\mu_i(u)$ and capacities $\rho_i(u)$ by

$$\mu_i(u) = \mu'_i(u)/\mu'_i(V) \quad \text{and} \quad \rho_i = d/\mu'_i(V).$$

We obtain an instance $\mathcal{I}'$. Note that the optimal solution $P_1^*, \ldots, P_k^*$ for $\mathcal{I}$ is a feasible solution for $\mathcal{I}'$ since

$$\begin{aligned}\mu_i(P_i^*) &= \sum_{u \in P_i^*} \frac{\mu'_i(u)}{\mu'_i(V)} = \frac{1}{\mu'_i(V)} \sum_{u \in P_i^*} \max_j \frac{r_j(u,i)}{c_j(i)} \leq \frac{1}{\mu'_i(V)} \sum_{u \in P_i^*} \sum_{j=1}^{d} \frac{r_j(u,i)}{c_j(i)} \\ &= \frac{1}{\mu'_i(V)} \sum_{j=1}^{d} \sum_{u \in P_i^*} \frac{r_j(u,i)}{c_j(i)} \leq \frac{d}{\mu'_i(V)} = \rho_i.\end{aligned}$$

We solve instance $\mathcal{I}'$ and get a partitioning $P_1, \ldots, P_k$ that cuts at most $O(\sqrt{\log n \log k}\, OPT)$ edges. The partitioning satisfies d-dimensional capacity constraints:

$$\sum_{u \in P_i} r_j(u,i) \leq \sum_{u \in P_i} c_j(i)\mu'_i(u) = c_j(i)\mu'_i(V) \sum_{u \in P_i} \mu_i(u) \leq \; \leq c_j(i)\mu'_i(V)(5(1+\varepsilon)\rho_i) = 5d(1+\varepsilon)\, c_j(i).$$

This concludes the analysis of the reduction.

Remark C.1. *The algorithm $\mathcal{A}$ from Theorem 2.1 is a randomized algorithm: it always finds a feasible solution (a solution with $|P_i| \leq 5(1+\varepsilon)\rho_i$), the expected cost of the solution is $\alpha_{\mathcal{A}} SDP = O(D \times OPT)$ (where $\alpha_{\mathcal{A}} = O(D)$), and the expected number of iterations the algorithm performs is upper bounded by $4n^2k + 1$. The algorithm can be easily converted to an algorithm $\mathcal{A}'$ that always runs in polynomial-time and that succeeds with high probability. If it succeeds, it outputs a feasible solution of cost $O(D \times OPT)$; if it fails, it outputs $\perp$ ($\perp$ is a special symbol that indicates that the algorithm failed). The algorithm $\mathcal{A}'$ works as follows. It executes $\mathcal{A}$. If $\mathcal{A}$ does not stop after $(4n^4k + n^2)$ iterations, $\mathcal{A}'$ terminates and outputs $\perp$. Otherwise, it compares the value of the solution that $\mathcal{A}$ found with $3\alpha_{\mathcal{A}}\, SDP$: If the cost is less than $3\alpha_{\mathcal{A}}\, SDP$, the algorithm outputs the solution; otherwise it outputs $\perp$. Clearly the algorithm always runs in polynomial time, and if it succeeds it finds a solution of cost at most $3\alpha_{\mathcal{A}}\, OPT = O(D \times OPT)$. By Markov's inequality, the probability that the algorithm fails is at most $1/n^2 + 1/3 < 1/2$. By running the algorithm n times, we can make the failure probability exponentially small (note that we need the algorithm to succeed at least once).*